

DEVICES AND METHODS FOR MINIMIZING START UP DELAY
IN TRANSMISSION OF STREAMING MEDIA

BACKGROUND OF THE INVENTION

1. Field of the invention

The present invention is related to the field of data transmission over networks, and more specifically to devices and methods for minimizing the startup delay of streaming media.

2. Description of the related art

Networks are increasingly used to transmit data. The data can be for various media, such as audio or video, and is usually organized in groups that are called frames. When such data is transmitted continuously, it is also called streaming media. The transmission can either be for live, or for canned media. Either way, the transmission has to be in real-time, or in delayed real-time.

A fundamental requirement for streaming media is that it has to be played at a constant rate. Packet networks, however, may transmit data asynchronously. Therefore, for streaming media to be transmitted over packet networks, strict timing requirements have to be met in order to achieve high-quality media play out. Furthermore, where there are two related but separate streams of media (such as audio and video), they have to be time-synchronized.

The bandwidth of the channel is intended to be wider than what is required by the streaming media. The channel bandwidth, however, is occasionally subject to congestion. This causes play out of streaming media to exhibit the problem of network jitter, if played out in real-time. This means that what is played out might be momentarily discontinued, if the bandwidth of the network is momentarily constricted to less than the required bandwidth of the streaming media.

Referring to Fig. 1, the problem of network jitter has been solved in the art by introducing a de-jitter buffer at the receiving end. A server or proxy 110 transmits data to a client or receiver 160 through a network 150. The server 110 includes a path 112 for data

5 received from a source, and transmitted to the network 150. Typically, the server 110 includes a buffer in the path 112, such as a transmit buffer 114.

Importantly, the client 160 includes a fixed size de-jitter receive buffer 162. The buffer 162 first fills up to its size, and then starts playing out. While playing out, the buffer 162 is emptied at the same rate as it is filled. Constrictions in bandwidth affect only the
10 fullness of the buffer, not the play out, therefore preventing network jitter. The larger the size of the buffer, the more the network jitter phenomenon is abated.

A problem generated by the de-jitter receive buffer 162 is a phenomenon called startup delay. The phenomenon is now described in more detail below.

Referring to Fig. 2, the fullness of the buffer 162 is plotted against a time axis as a
15 line 210. Initially the buffer 162 fills up to its size f_y , which it reaches at time point 212. For illustration purposes, the line 210 is shown in an approximate manner as a ramp. However, the line is actually piecewise linear, shaped more like a stair.

After reaching its full size f_y , playout begins. Then the buffer 162 ideally remains at the same fullness level. Again, that is shown as a straight line only for illustration purposes.
20 The fullness level buffer 162 may temporarily go lower than the ideal.

It will be appreciated that the larger the size f_y is, the longer time it takes to reach that level. The time between zero and time point 212 is called the startup delay. The startup delay causes the play out to be in delayed real-time, compared to when the data is received.

The startup delay can be as long as one second or more, which is a problem. It is
25 even more of a problem when the user is channel surfing, i.e. switching quickly from one channel to another. With every new channel switch, there is another startup delay.

BRIEF SUMMARY OF THE INVENTION

The present invention overcomes the problems and limitations of the prior art.

30 Generally, the present invention provides devices and methods that diminish the startup delay of streaming media.

In one embodiment, the invention maintains a portion of the media stream stored in an initial burst transmit buffer of the server. At startup, the stored portion is transmitted at a rate higher than the fixed frame rate. This exploits the additional, unused bandwidth of the
35 network. This initial transmission fills up the de-jitter receive buffer at the receiving end more quickly, thereby shortening the startup delay. Then transmission is switched to the regular rate.

5 In addition, a variable bit rate transcoder is optionally used for the data of the initial transmission. The transcoder diminishes somewhat the size of these frames, so they can be transmitted faster. This way the start up delay is shortened even more.

Moreover, the fill level of the de-jitter receive buffer may be started at a value lower than a final value. This speeds up the beginning of play out, further shortening the delay
10 time. The fill level of the de-jitter receive buffer is then gradually increased to a desired final value.

The invention will become more readily apparent from the following Detailed Description, which proceeds with reference to the drawings, in which:

15 BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 is a block diagram of a server transmitting data to a client over a network.

Fig. 2 is a time diagram that demonstrates why the phenomenon of a startup delay occurs in the arrangement of Fig. 1.

Fig. 3 is a block diagram of a server made according to the invention transmitting data
20 over a network to a client made according to the invention.

Fig. 4 is a flowchart illustrating a first method according to the invention.

Fig. 5 is a time diagram comparing the improved startup delays accomplished by various embodiments of the invention to the startup delay of Fig. 2.

Fig. 6 is a flowchart illustrating a second method according to the invention.

Fig. 7 is a time diagram illustrating the changing fill level of a de-jitter buffer
25 according to the invention.

Fig. 8 is a diagram illustrating how the improved start up delays afforded by the present invention depend on a ratio of a demanded bandwidth over the available network transmission bandwidth.

30 DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT(S)

As has been mentioned, the present invention provides devices and methods for minimizing the startup delay of streaming media. Various aspects and embodiments of the invention are now described in more detail. The system of the invention can be a server, a
35 client, or a combination of them.

Referring to Fig. 3, a device according to the invention is a server 310, which is also known as a proxy 310. As shown, the server 310 can be either a full server, or a portion of a server. The server 310 is for transmitting data over a network 150, such as a packet network,

5 to a client 360. The client 360, which is also known as a receiver 360, includes at least one de-jitter buffer.

The server 310 includes a regular path 312 for transmitting to the network 150 data received from a source. Preferably, the transmit buffer 314 is located in the regular path 312, for buffering data from the source. Transmission through the regular path 312 is intended to
10 be at the regular rate, such as at a fixed frame rate r required by the streaming media in question. The frame rate r may be, for example, 100 frames per second for all audio, or 25 frames per second for video. The streaming media may have a constant or a variable transmission rate.

The server 310 also includes a burst path 322. The burst path 322 is for transmitting
15 data received from the source at a burst rate. The burst rate is higher than the regular rate.

The server 310 most preferably includes an initial burst transmit buffer 324 in the burst path 322. The initial burst transmit buffer 324 is for buffering the data received from the source, above and beyond the operations of transmit buffer 314. The initial burst transmit buffer 324 is also for transmitting the buffer to data through the network 150 to the client 360
20 at the burst rate.

The server 310 also includes a switch 330. The switch 330 is for selecting to transmit data to the network 150 from either the regular path 312, or the burst path 322. After the switch, the regular path 312 and the burst path 322 are merged.

The server 310 preferably also comprises a control unit 335. The control unit 335 is
25 primarily for controlling the switch 330. The control unit 335 can be implemented as a processor, software in a processor, a separate circuit, etc.

The server 310 preferably also includes a monitor that measures an amount of the data that is output through the burst path 322. That measure is most conveniently a number of frames, the same number as are stored in initial burst transmit buffer 324. This monitor can
30 be implemented in connection with the initial burst transmit buffer 324, such as by software. Alternately, this monitor can be implemented in connection with the control unit 335. The control unit 335 switches the switch 330, when a preset measure of the data is output through the burst path 322.

The server 310 moreover includes a bandwidth monitor that can measure the
35 bandwidth of the network 150. This feature can be optionally implemented in conjunction with control unit 335. This feature is used to provide an estimate of how fast data is being received at the client 360.

5 Additionally, the server 310 can include a fill level controller that controls the fill level of the initial burst transmit buffer 324. The fill level can be controlled according to the monitored bandwidth.

 The server 310 preferably also comprises a transcoder 340. The transcoder 340 is made as is known in the art, and is for transcoding the buffered streaming media that this
10 output through the burst path 322. The transcoder 340 could well be placed exclusively in the burst path 322. It is preferred, however, to place the transcoder 340 where the regular path 312 and the burst path 322 are merged. This way, the transcoder 340 can be additionally used for transcoding the subsequent regular transmission, if it is so desired. Optionally and preferably, the transcoder 340 transcodes the data of the initial burst transmit buffer 324, if
15 the monitored bandwidth of the network 150 becomes less than a preset bandwidth.

 A client 360 according to the invention receives streaming media from the network 150. In other words, the client 360 receives the data that encodes the streaming media. The client 360 can be any device that receives streaming media from a network, such as a network appliance.

20 The client 360 has a receive de-jitter buffer 364 for receiving the streaming media. In addition, it includes a variable bit rate decoder 366. It is preferred that the receive de-jitter buffer 364 has a changing fill level. It is preferred that the fill level is increased gradually, as is described later. It should be noted, however, that the server 310 of invention does not need the client 360 to have a variable size de-jitter buffer.

25 Referring to Fig. 4, a method of the invention is described with reference to flowchart 400. The method is for a server to retransmit streaming media to a network. The streaming media is encoded in data, and ultimately destined for a client.

 According to an optional box 410, the available bandwidth of transmission is monitored. This can be accomplished by monitoring the available bandwidth of the network.

30 According to an optional box 420, a size is adjusted of a first portion of streaming media. This size can be the number of frames that are to be stored in an initial transmit buffer, such as buffer 324.

 According to a box 430, a first portion of the streaming media is received along a burst path of the server. The size of the first portion may have been adjusted as per the
35 above. Preferably, the first portion is continuously updated, such as for live media.

 According to an optional box 440, the streaming media received along the burst path is stored in an initial burst transmit buffer, such as buffer 324.

5 According to an optional box 450, it is inquired whether the monitored bandwidth is less than a preset bandwidth. If it is, then according to box 455, the first portion of the streaming media is transcoded. It can be transcoded as it is being played out.

According to a box 460, the first portion of the streaming media (whether transcoded or not) is transmitted to the network through the burst path at the burst rate. The burst rate is
10 as high as permitted by the bandwidth of the connection.

According to a box 470, a second portion of the streaming media is received from the source. The second portion is received along a regular path of the server, which is distinct from the burst path at least in part.

According to a box 480, the server is switched to output data from the regular path,
15 instead of the burst path.

According to a box 490, the second portion is transmitted to the network through the regular path. It is transmitted at a regular rate, which is lower than the burst rate. The regular rate is intended to be the normal frame rate of the streaming media.

Advantages of the invention can be now appreciated with reference to Fig. 5. The
20 fullness of the buffer 364 is plotted in a vertical axis 504, against a time axis 510. Figure 5 is not necessarily to scale.

The performance of the prior art systems, which was explained in detail with reference to Fig. 2, is presented again in Fig. 5. That was line 210, which gave a start-up time at time point 212.

25 The initial burst transmit buffer 324 transmits faster than the regular rate of line 210. The resulting performance is given by line 326, which has a higher rate of ascent, because it is using the full bandwidth of the network 150. Line 326 levels out after the buffer 364 reaches its full level fy. This happens at time point 328, at which play out is triggered. This occurs earlier than time point 212, which corresponds to a shortening of the start-up delay.
30 Again, line 326 is straight only ideally.

Moreover, transcoding the first portion further accelerates how fast the buffer 364 reaches its nominal fullness level fy. That is because the fullness is measured in terms of frames, and the transcoded first portion includes less data per frame. This additional feature is depicted by line 342, which has an even higher rate of ascent than line 326. Accordingly,
35 line 342 reaches the fullness level fy at time point 344, which is even earlier than time point 328, thus further shortening the start-up delay. When this feature is implemented, it is true that the first few frames are of a lower quality, because of the transcoding. This does not

5 matter that much, however, because that is also the time that the end user is getting used to the new stream of information.

Referring now to Fig. 6, a second method of invention is described with reference to flowchart 600. The second method is for client to receive streaming media over a network.

According to a box 610, the client receives data that encode the streaming media from
10 the network.

According to a box 620, the received data is stored in a de-jitter buffer. This increases a fullness of the buffer.

According to a box 630, it is inquired whether the fullness has reached a first fill level. If it has, then according to a box 640, play out of the stored data is initiated.

15 According to a box 650, the fill level of the de-jitter buffer is changed while playing out the stored data. Preferably, it is increased gradually up to a final fill level. In other words, play out is initiated before the fullness has reached the final fill level.

Referring to Fig. 7, the fill level of buffer 364 is plotted against time by line 710. Up until time point 712, the fill level equals fx . The fill level, or buffer size, of buffer 364 is
20 gradually increased, preferably linearly. It reaches a final fill level fy at time point 716.

The advantage of changing gradually the fill level of the de-jitter buffer will be appreciated by referring back to Fig. 5. Assuming that both the initial burst transmit buffer 324 and transcoder 340 are enabled, the fullness of buffer 364 will be given by line 714. This starts identically from the origin 0, until he teaches a level fx at time point 712. This initiates
25 play out, at the time even earlier than time point 344. This corresponds to the yet fastest start-up delay of all the contrasted schemes.

After time point 712, as the fill level is increased, so is the fullness level. As seen before, the fill level will reach level fy at time point 716. The fullness would also reach fy , at a time point 718. In general, time point 718 occurs later than time point 716.

30 It should be noted that both time points 716 and 718 occur after time point 212, but that does not matter. The play out has been initiated at time point 712, which shortens the start-up delay.

It should be further noted that the fill level fx started increasing at time point 712, i.e. when the initial fill level fx was reached. That is not necessary, however. The fill level of
35 buffer 364 can start increasing to its final level before or after time point 712.

The invention offers its highest advantage when a high bandwidth is available for the initial, burst transmission. This is now quantified for the various embodiments discussed above.

5 Referring now to Fig. 8, the start-up delay is plotted against a ratio of BS/BI. BS is the transmission rate required by the streaming media. BI is the actually available bandwidth. For values larger than 1, no sustained streaming media transmission is meaningful.

The start-up delay of the prior art is given by line 810. Line 810 corresponds to the performance of the prior art system, discussed above with reference to Fig. 2. In other words,
10 the start-up delay is the same, regardless of how much available bandwidth BI there is over the required bandwidth BS.

Once the ratio of BS/BI is known, a ratio point 806 is chosen on the ratio axis. Then a vertical line 510 is drawn from ratio point 806. It should be noted that this line 510 is also the horizontal line 510 in the diagram of Fig. 5. Accordingly, where line 810 intersects line
15 510, time point 212 is defined.

If only the initial burst transmit buffer 324 is implemented, the behavior of the system is characterized by line 830. Line 830 intersects line 510 at time point 328.

If also the transcoder 340 transcodes the data stored in buffer 324, then the behavior of the system is characterized by line 840. Line 840 intersects line 510 at time point 344.

Sub A: If, in addition to the above, the buffer fill level is increased gradually from f_x to f_y , then the behavior of the system is characterized by line 870. Line 870 intersects line 510 at time point 712.

As can be seen from the vertical axis of Fig. 8, the start-up delay becomes smaller, as more and more features of the invention are implemented. In addition, as more bandwidth is
25 available, the line 510 shifts closer to the left. Then the time points 712, 344, and 328 slide down accordingly, corresponding to a diminishing startup delay. In other words, the more bandwidth becomes available, the better the invention works in minimizing the start-up delay.

A person skilled in the art will be able to practice the present invention in view of the description present in this document, which is to be taken as a whole. Numerous details have
30 been set forth in order to provide a more thorough understanding of the invention. In other instances, well-known features have not been described in detail in order not to obscure unnecessarily the invention.

While the invention has been disclosed in its preferred form, the specific embodiments thereof as disclosed and illustrated herein are not to be considered in a limiting
35 sense. Indeed, it should be readily apparent to those skilled in the art in view of the present description that the invention can be modified in numerous ways. The inventor regards the subject matter of the invention to include all combinations and subcombinations of the various elements, features, functions and/or properties disclosed herein.

5 The following claims define certain combinations and subcombinations, which are regarded as patentable. Additional claims for other combinations and subcombinations of features, functions, elements and/or properties may be presented in this or a related document.